

Why Just Boogie?

Translating Between Intermediate Verification Languages

Michael Ameri¹ and Carlo A. Furia^{2*}

¹ Chair of Software Engineering, Department of Computer Science,
ETH Zurich, Switzerland mameri@student.ethz.ch

² Department of Computer Science and Engineering,
Chalmers University of Technology, Sweden furia@chalmers.se

Abstract. The verification systems Boogie and Why3 use their respective intermediate languages to generate verification conditions from high-level programs. Since the two systems support different back-end provers (such as Z3 and Alt-Ergo) and are used to encode different high-level languages (such as C# and Java), being able to translate between their intermediate languages would provide a way to reuse one system’s features to verify programs meant for the other. This paper describes a translation of Boogie into WhyML (Why3’s intermediate language) that preserves semantics, verifiability, and program structure to a large degree. We implemented the translation as a tool and applied it to 194 Boogie-verified programs of various sources and sizes; Why3 verified 83% of the translated programs with the same outcome as Boogie. These results indicate that the translation is often effective and practically applicable.

1 Introduction

Intermediate verification languages (IVLs) are intermediate representations used in verification technology. Just like compiler design has benefited from decoupling front-end and back-end, IVLs help write verifiers that are more modular: the front-end specializes in encoding the rich semantics of a high-level language (say, an object-oriented language such as C#) as a program in the IVL; the back-end generates verification conditions (VCs) from IVL programs in a form that caters to the peculiarities of a specific theorem prover (such as an SMT solver).

Boogie [2] and WhyML [8] are prime examples of popular IVLs with different, often complementary, features and supporting systems (respectively called Boogie and Why3). In this paper we describe a translation of Boogie programs into WhyML programs and its implementation as the tool `b2w`. As we illustrate with examples in Sec. 3, using `b2w` increases the versatility brought by IVLs: without having to design and implement a direct encoding into WhyML or even being familiar with its peculiarities, users can take advantage of some of the best features of Why3 when working with high-level languages that translate to Boogie.

* Work done mainly while affiliated with ETH Zurich.

Boogie vs. WhyML. While the roles of Boogie and WhyML as IVLs are similar, the two languages have different characteristics that reflect a focus on complementary challenges in automated verification. Boogie is the more popular language in terms of front-ends that use it as IVL, which makes a translation *from* Boogie more practically useful than one into it; it has a finely tuned integration with the Z3 prover that results from the two tools having been developed by the same group (Microsoft Research’s RiSE); it combines a simple imperative language with an expressive typed logic, which is especially handy for encoding object-oriented or, more generally, heap-based imperative languages. In contrast, WhyML has a more flexible support for multiple back-end provers it translates to, including a variety of SMT solvers as well as interactive provers such as Coq; it can split VCs into independent goals and dispatch them to different provers; it offers limited imperative constructs within a functional language that belongs to the ML family, which brings the side benefit of being able to *execute* WhyML programs—a feature quite useful to debug and validate verification attempts.

Goals and evaluation. The overall goal of this paper is devising a translation \mathcal{T} from Boogie to WhyML programs. The translation, described in Sec. 4, should preserve correctness, verifiability, and readability as much as possible. Preserving correctness means that, given a Boogie program p , if its translation $\mathcal{T}(p)$ is a correct WhyML program then p is correct (soundness); the converse should also hold as much as possible: if $\mathcal{T}(p)$ is incorrect then p is too (precision). Preserving verifiability means that, given a Boogie program p that verifies in Boogie, its translation $\mathcal{T}(p)$ is a WhyML program that verifies in Why3. Preserving readability means that the translation should not introduce unnecessary changes in the structure of programs.

The differences, outlined above, between Boogie and WhyML and their supporting systems make achieving correctness, verifiability, and readability challenging. While we devised \mathcal{T} to cover the entire Boogie language, its current implementation *b2w* does not fully support a limited number of features (branching, the most complex polymorphic features, and bitvectors) that make it hard to achieve verifiability in practice. In fact, while replacing branching (*goto*) with looping is always possible [11], a general translation scheme does not produce verifiable loops since one should also infer invariants (which are often cumbersome due to the transformation). Polymorphic maps are supported to the extent that their type parameters can be instantiated with concrete types; this is necessary since WhyML’s parametric polymorphism cannot directly express all usages in Boogie, but it may also introduce a combinatorial explosion in the translation; hence, *b2w* fails on the most complex instances that would be unmanageable in Why3. Boogie’s bitvector support is much more flexible than what provided by Why3’s libraries; hence *b2w* may render the semantics of bitvector operations incorrectly.

These current implementation limitations notwithstanding (see Sec. 4 for details), we experimentally demonstrate that *b2w* is applicable and useful in practice. As Sec. 5 discusses, we applied *b2w* to 194 Boogie programs of different size and sources; most of the programs have not been written by us and exercise Boogie in a variety of different ways. For 83% (161) of these programs, *b2w* produces a WhyML translation that Why3 can verify as well as Boogie can verify the original, thus showing the feasibility of automating translation between IVLs.

Tool availability. The tool b2w is available as open source at:

https://bitbucket.org/michael_ameri/b2w/

2 Related Work

Translations and abstraction levels. Translation is a ubiquitous technique in computer science; however, the most common translation schemes bridge *different abstraction levels*, typically encoding a program written in a high-level language (such as Java) into a lower-level representation which is suitable for execution (such as byte or machine code). Reverse-engineering goes the opposite direction—from lower to higher level—for example to extract modular and structural information from C programs and encode it using object-oriented constructs [26]. This paper describes a translation between intermediate languages—Boogie and Why3—which belong to *similar abstraction levels*. In the context of model transformations [20], so-called bidirectional transformations [25] also target lossless transformations between notations at the same level of abstraction.

Intermediate verification languages. The Spec# project [3] introduced Boogie to add flexibility to the translation between an object-oriented language (a dialect of C#) and the verification conditions in the logic fragments supported by SMT solvers. An intermediate verification language embodies the idea of intermediate representation—a technique widespread in compiler construction—in the context of verification. Since its introduction for Spec#, Boogie has been adopted as intermediate verification language for numerous other front-ends such as Dafny [16], AutoProof [28], Viper [12], and Joogie [1]; its popularity demonstrates the advantages of using intermediate verification languages.

While Boogie retains some support for different back-end SMT solvers, Z3 [7] remains its fully supported primary target. By contrast, supporting multiple, different back-ends is one of the main design goals behind the Why3 system [8], which does not merely generate verification conditions in different formats but offers techniques to split them into independently verifiable units and to dispatch each unit to a different prover. Why3 also fully supports interactive provers,³ which provide a powerful means of discharging the most complex verification conditions that defy complete automation.

Another element that differentiates Boogie and Why3 is the support for executing programs; this is quite useful for debugging verification attempts and for applying testing-like techniques to the realm of verification. Boogaloo [21] supports symbolic execution of Boogie programs; Symbooglix is a more recent project with the same goal [19]. Thanks to it being a member of the ML family, Why3 directly supports symbolic execution as well as compilation of WhyML programs to OCaml.

In all, while the Boogie and WhyML languages belong to a similar abstraction level, they are part of systems with complementary features, which motivates this paper’s idea of translating one language into the other. Since Boogie is overall more popular, in terms of tools that use it as a back-end, the translation from Boogie to WhyML is more practically useful than the one in the opposite direction.

³ In comparison, Boogie’s support for HOL is restricted and not up-to-date [4].

Other intermediate languages for verification are Pilar [24], used in the Sireum framework for SPARK; Silver [12], an intermediate language with native support for permissions in the style of separation logic; and the flavor of dynamic logic for object-oriented languages [23] used in the KeY system. Another approach to generalizing and reusing different translations uses notions from model transformations to provide validated mappings for different high-level languages [5]. Future work may consider supporting some of these intermediate languages and approaches.

3 Motivating Examples

Verification technology has made great strides in the last decade or two, but a few dark corners remain where automated reasoning shows its practical limitations. Fig. 1 provides three examples of simple Boogie programs that trigger incorrect or otherwise unsatisfactory behavior. We argue that translating these programs to WhyML makes it possible to verify them using a different, somewhat complementary verification tool; overall, confidence in the results of verification is improved.

Procedure `not_verify` in Fig. 1 has a contradictory postcondition (notice $N < N$, N is a nonnegative constant, and the loop immediately terminates). Nonetheless, recent versions of Boogie and Z3 successfully verify it.⁴ More generally, unless the complete toolchain has been formally verified (a monumental effort that has only been performed in few case studies [18,13,14]), there is the need to *validate* the successful runs of a verifier. Translating Boogie to Why3 provides an effective validation, since Why3 has been developed independent of Boogie and uses a variety of backends that Boogie does not support. Procedure `not_verify` translated to Why3 (Fig. 2) does not verify as it should.

Procedures `lemma_yes` and `lemma_no` in Fig. 1 demonstrate Boogie’s support for mathematical real numbers, which is limited in the way the power operator `**` is handled. Boogie vacuously verifies both properties $2^3 > 0$ and $2^3 < 0$, even though Z3 outputs some unfiltered errors that suggest the verification is spurious (the power operator `**` is not properly supported); indeed, only the inequality encoded by `lemma_yes` is correct. Why3 provides a more thorough support for real arithmetic, both by translating to backends such as Alt-Ergo and by providing a more effective encoding in Z3; in fact, it verifies the translated procedure `lemma_yes` but correctly fails to verify `lemma_no`.

The loop in procedure `trivial_inv` in Fig. 1 includes an invariant asserting that `i` takes only even values. Even if this is clearly true, Boogie fails to check it; pinning down the precise cause of this shortcoming requires knowledge of Boogie’s (and Z3’s) internals, although it likely is a manifestation of the “triggers” heuristics that handle (generally undecidable) quantified expressions. Based on this knowledge, there are specification patterns that try to work around such idiosyncrasies; in the example, one could introduce a “witness” ghost variable `k` such that $i = 2*k$ is an invariant. However, if we insist on verifying the program in its original form, Why3 can dispatch verification conditions to *interactive* provers, where the user provides the crucial proof steps.⁵ Cases such as the loop invariant of `trivial_inv` where a proof is “obvious” to a human user

⁴ <https://github.com/boogie-org/boogie/issues/25>

⁵ Why3 can also check the invariant automatically by relying on the CVC4 SMT solver.

but it clashes against the default strategies to handle quantifiers are prime candidate to exploit interactive provers. Thus, translating Boogie to Why3 provides another means of exploiting the latter’s versatile support for interactive provers and multiple backends.

```

const N: int;
axiom 0 ≤ N;

procedure not_verify()
  ensures (∀ k, l: int •
    0 ≤ k ≤ l < N ⇒ N < N);
{
  var x: int;
  x := -N;
  while (x ≠ x) { }
}

procedure lemma_yes()
  ensures 2.0**3.0 > 0.0;
{ }

procedure lemma_no()
  ensures 2.0**3.0 < 0.0;
{ }

procedure trivial_inv()
{
  var i: int;
  i := 0;
  while (i < 10)
    invariant 0 ≤ i ≤ 10;
    invariant
      (∃ j: int • i = 2*j);
    { i := i + 2; }
}

```

Fig. 1. Three simple Boogie programs for which automated reasoning is limited.

4 Boogie-to-Why3 Translation

Intermediate languages for verification combine programming constructs and a logic language. When used to encode programs written in a high-level language, the programming constructs encode program behavior, and the logic constructs encode specifications, constrain the semantics to conform to the high-level language’s (typically through axioms), and support other kinds of annotations (such as triggers).

Both Boogie and WhyML provide, as logic language, a typed first-order logic with arithmetic. Boogie’s programming constructs are a simple imperative language with both structured (while loops, procedures) and unstructured (jumps, global variables) statements. WhyML’s programming constructs combine an ML-like functional language with a few structured imperative features such as mutable variables and loops.

Correspondingly, we define a translation $\mathcal{T}: \text{Boogie} \rightarrow \text{WhyML}$ of Boogie to WhyML as the composition $\mathcal{E} \circ \mathcal{D}$ of two translations: $\mathcal{D}: \text{Boogie} \rightarrow \text{Boogie}$ is a desugaring⁶ which rewrites away the Boogie constructs, such as *call-forall*, that have no similar construct in WhyML by expressing them using other features of Boogie. Then, $\mathcal{E}: \text{Boogie} \rightarrow \text{WhyML}$ encodes Boogie programs simplified by \mathcal{D} as WhyML programs, while introducing constraints that ensure that the semantics in WhyML mirrors the one in Boogie. For simplicity, the presentation does not sharply separate the two translations \mathcal{D} and \mathcal{E} but defines either or both of them as needed to describe the translation of arbitrary Boogie constructs.

A single feature of the Boogie language significantly compounds the complexity of the translation: *polymorphic maps*, which correspond to mappings between domains of generic type. Why3 does support polymorphic maps through a library, but its type system is more restrictive and does not allow the same degree of freedom as Boogie’s

⁶ This is unrelated to Boogie’s built-in desugaring mechanism (option `/printDesugared`).

```

constant N: int
axiom A0: 0 ≤ N;

val not_verify (): ()
  ensures { ∀ k, l: int .
    0 ≤ k ≤ l < N → N < N }

let not_verify_impl(): ()
  ensures { ∀ k, l: int .
    0 ≤ k ≤ l < N → N < N }
  =(
    let x = ref (any int) in
    x.contents ← -N;
    while
      (x.contents ≠ x.contents)
    do done;
  end )

val lemma_yes (): ()
  ensures
    { (pow 2.0 3.0) >. 0.0 }

val lemma_no (): ()
  ensures
    { (pow 2.0 3.0) <. 0.0 }

let lemma_yes_impl (): ()
  ensures
    { (pow 2.0 3.0) >. 0.0 }
  =( )

let lemma_no_impl (): ()
  ensures
    { (pow 2.0 3.0) <. 0.0 }
  =( )

val trivial_inv (): ()

let trivial_inv_impl (): ()
  =(
    let i = ref (any int) in
    i.contents ← 0;
    while (i.contents < 10) do
      invariant
        { 0 ≤ i.contents ≤ 10 }
      invariant
        { ∃ j: int .
          i.contents = 2*j }
      i.contents ← i.contents + 2;
    done;
  )

```

Fig. 2. The translation to WhyML of the three Boogie programs in Fig. 1. (Boilerplate such as general declarations, imports, and frame condition checking are omitted for clarity.)

in using variables of polymorphic map types. For clarity, the presentation of the translation initially ignores polymorphic maps. Then, Sec. 4.10 discusses how the general translation scheme can be extended to support them.

As running examples, Fig. 2 shows how \mathcal{T} translates the examples of Fig. 1.

4.1 Types

Boogie types include primitive types, instantiated type constructors, and map types.

Primitive types are **int** (mathematical integers), **real** (mathematical reals), **bool** (Booleans), and **bv** n (n -bit vectors). \mathcal{T} translates primitive types into their Why3 analogues as shown in Tab. 3. Since Why3 offers primitive types and operations on them through libraries, \mathcal{T} also generates import statements for the libraries that provide the same operations that are available in Boogie, such as integer to/from real conversion.

T	$\mathcal{T}(T)$	Why3 libraries
int	int	int.Int, int.EuclideanDivision
real	real	real.RealInfix, real.FromInt, real.Truncate, real.PowerReal
bool	bool	bool.Bool
bv n	bv	bv.BitVector with constant size = n

Table 3. Translation of primitive types, and Why3 libraries supplying the necessary operations.

Type constructors. A Boogie type declaration using the *type constructor* syntax⁷ introduces a new parametric type T with parameters a_1, \dots, a_m . \mathcal{T} translates it to an alge-

⁷ \mathcal{T} ignores the optional type modifier **finite**, since it does not seem fully supported in Boogie.

braic type with constructor T : $\mathcal{T}(\text{type } T \ a_1 \dots a_m) = \text{type } T \ 'a_1 \dots 'a_m$ for $m \geq 0$, where ticks ' identify type parameters in WhyML.

Map types. A Boogie *map type* M declared as: $\text{type } M = [T_1, \dots, T_n] \ U$ defines the type of a mapping from $T_1 \times \dots \times T_n$ to U , for $n \geq 1$. Why3 supports maps through its library `map.Map`,⁸ hence, $\mathcal{T}(M) = \text{map } (\mathcal{T}(T_1), \dots, \mathcal{T}(T_n)) \ \mathcal{T}(U)$, where an n -tuple encapsulates the n -type domain of M .

4.2 Constants

The translation of constant declarations is generally straightforward, following the scheme:

$$\mathcal{T}(\text{const } c : T) = \text{constant } c : \mathcal{T}(T)$$

Unique constants. All constants of a type T declared with the modifier `unique` have values that are pairwise different. Thus, for m constants `const unique` $c_1, \dots, c_m : T$, \mathcal{T} encodes the uniqueness properties using $\binom{m}{2}$ axioms `axiom unique_c_i_j`: $c_i \neq c_j$, for $1 \leq i \neq j \leq m$.

Orders. Boogie provides the operator $<$: to express partial order over every type; \mathcal{T} introduces a polymorphic operator $<$: and axiomatizes its reflexive, antisymmetric, and transitive properties:

```
predicate (<:) (x: 'a) (y: 'a)
axiom ReflexiveP0:  $\forall x : 'a. x <: x$ 
axiom AntisymmetricP0:  $\forall x y : 'a. x <: y \wedge y <: x \rightarrow x = y$ 
axiom TransitiveP0:  $\forall x y z : 'a. x <: y \wedge y <: z \rightarrow x <: z$ 
```

Boogie supplies special syntax to describe a partial-order relations with a certain structure, which corresponds to a DAG where any two nodes x and y are connected by an edge $x \rightarrow y$ iff $x <: y$ and y is a direct successor of x in the order. Let a, b, c, d, e, f be unique⁹ constants of the same type T . The Boogie syntax to specify ordering between them is in Fig. 4. \mathcal{D} reconstructs the DAG of the order specification, and then formalizes it in axiomatic form. For example, the specifications in Fig. 4 determine the DAG in Fig. 5, which is axiomatized as in Fig. 6.

4.3 Variables

Why3 supports mutable variables through the reference type `ref` from theory `Ref`. Boogie global variable declarations become global value declarations of type `ref`; Boogie local variable declarations become `let` bindings with local scope. Thus, if v is a global variable and $_v$ is a local variable in Boogie:

```
global variable  $\mathcal{T}(\text{var } v : T) = \text{val } v : \text{ref } \mathcal{T}(T)$ 
local variable  $\mathcal{T}(\text{var } \_v : T) = \text{let } \_v = \text{ref } (\text{any } \mathcal{T}(T)) \text{ in}$ 
```

The expression `any` T provides a nondeterministic value of type T .

⁸ Why3's maps, like Boogie's, do not satisfy extensionality (<http://lists.gforge.inria.fr/pipermail/why3-club/2013-February/000572.html>).

⁹ Uniqueness is not required but makes the order specification easier to present.

BOOGIE SPECIFICATION	SEMANTICS
<code>const c: T extends a, b;</code>	a and b are the only direct successors of c
<code>const a: T extends;</code>	a has no (direct) successors
<code>const d: T extends c complete;</code>	c has no direct predecessors other than d and any others that are explicitly specified
<code>const e, f: T extends unique d;</code>	d is the only direct predecessor of both e and f, and the sub-graphs that originate in e and f are disjoint

Fig. 4. Ordering specifications in Boogie (older versions of Boogie use $<$: instead of **extends**).




Fig. 5. DAG corresponding to the ordering specification of Fig. 4. Solid edges denote the successor relation; dotted edges denote allowed (but not specified) relations; the dashed line expresses disjointness of the two sub-graphs.

4.4 Functions

Boogie function *declarations* become WhyML function declarations:

$$\begin{aligned} \mathcal{T}(\text{function } f(x_1: T_1, \dots, x_n: T_n) \text{ returns } (U)) \\ = \text{function } f(x_1: \mathcal{T}(T_1)) \cdots (x_n: \mathcal{T}(T_n)): \mathcal{T}(U) \end{aligned} \quad (1)$$

WhyML function *definitions* require, unlike Boogie's, a variant to ensure that recursion is well-formed. Therefore, Boogie function definitions are not translated into WhyML function definitions but are axiomatized: if function f in (1) has body B , \mathcal{D} replaces the body with the **axiom** $(\forall z_1: T_1, \dots, z_n: T_n \bullet f(z_1, \dots, z_n) = B)$.¹⁰

4.5 Expressions

The translation of Boogie expressions to WhyML expressions is mostly straightforward, given the translation of types described above. We describe the few cases that deserve some detail.

¹⁰ To take advantage of Why3's well-formedness checks, we plan to offer translations of Boogie functions to WhyML functions as a user option in future work.

```

axiom (c <: a  $\wedge$  c <: b  $\wedge$   $\forall x: T \bullet c <: x \implies c = x \vee a <: x \vee b <: x$ )
axiom ( $\forall x: T \bullet \neg(a <: x)$ )
axiom (d <: c  $\wedge$   $\forall x: T \bullet x <: c \implies c = x \vee x <: d$ )
axiom (e <: d  $\wedge$   $\forall x: T \bullet e <: x \implies e = x \vee d <: x$ )
axiom (f <: d  $\wedge$   $\forall x: T \bullet f <: x \implies f = x \vee d <: x$ )
axiom ( $\forall x: T \bullet x <: e \implies \neg(x <: f)$ )
axiom ( $\forall x: T \bullet x <: f \implies \neg(x <: e)$ )

```

Fig. 6. Axiomatization of the ordering specification in Fig. 4.

Nondeterministic choice. The special value `*` represent a nondeterministic Boolean choice (used in loop exit flags and conditionals); we define $\mathcal{T}(\ast) = \text{any } \text{bool}$, which provides a nondeterministic Boolean value.

Variables. Since a Boogie variable v of type T turns into a value v of type $\text{ref } \mathcal{T}(T)$, occurrences of v in an expression translate to $v.\text{contents}$, which represents the value attached to reference v .

Map expressions. \mathcal{T} translates map selection and update using functions `get` and `set` from theory `Map`. If m is a map of type M defined in Sec. 4.1, then:¹¹

	E	$\mathcal{T}(E)$
selection	$m[e_1, \dots, e_n]$	$\text{get } \mathcal{T}(m) (\mathcal{T}(e_1), \dots, \mathcal{T}(e_n))$
update	$m[e_1, \dots, e_n := f]$	$\text{set } \mathcal{T}(m) (\mathcal{T}(e_1), \dots, \mathcal{T}(e_n)) \mathcal{T}(f)$

Lambda expressions. Boogie recently introduced lambda expressions as syntactic sugar for maps. While WhyML has lambda abstractions, they are not allowed as first-order values in programs [6]. Instead, the translation desugars lambda expression into constant maps: $\mathcal{D}(\lambda x_1 : T_1, \dots, x_n : T_n \bullet e) = \text{lmb}$, where $\text{const lmb} : [T_1, \dots, T_n] \tau(e)$ is axiomatized by **axiom** $(\forall x_1 : T_1, \dots, x_n : T_n \bullet \text{lmb}[x_1, \dots, x_n] = e)$, and $\tau(e)$ is e 's type.

Old expression. Within a procedure's postcondition or body, the expression **old**(e) refers to the value of e in the prestate. WhyML offers a more general construct to refer to an expression's value at any labeled point within a procedure's body. Hence, every WhyML procedure implementation translating a Boogie procedure implementation includes a label "begin", so that $\mathcal{T}(\text{old}(e))$ is just **old** $\mathcal{T}(e)$ within postconditions, and is **at** $\mathcal{T}(e)$ ' "begin" within bodies.

Bitvectors. Why3's theory `BitVectors` does not provide all operations that are supported by Boogie. In particular, it does not support *extraction expressions* $b[n:m]$ (drop the m least significant bits and return the next $n - m$ least significant bits) and *concatenation expressions* $b ++ c$ (the bit vector obtained by concatenating b and c). \mathcal{T} introduces functions `extract` $(b: \text{bv}) (n: \text{int}) (m: \text{int}): \text{bv}$ and `cat` $(b: \text{bv}) (c: \text{bv}): \text{bv}$ and uses them to translate applications of these bit vector operators, but leaves them uninterpreted in Why3. \mathcal{T} 's implementation currently supports only the bitvectors operations available in Why3's theory `BitVectors`.

4.6 Procedures

Boogie procedures have a declaration (signature and specification) and zero or more implementations. The latter follow the general syntax of Fig. 7 (left), where a procedure p with input argument t and output argument u has one implementation with local variable l and body B . For simplicity of presentation, p has one input argument, one output argument, and one local variable, but generalizing the description to an arbitrary number of variables is straightforward.

¹¹ Despite its name, `set` returns a new map rather than changing its argument's value.

```

procedure p(t: T where Wt)
  returns (u: U where Wu);
  requires R;
  free requires fR;
  modifies M;
  ensures E;
  free ensures fE;

implementation p(t: T)
  returns (u: U)
  {
    var l: L where Wl;
    B
  }

```

```

val p (t :  $\mathcal{T}(T)$ ):  $\mathcal{T}(U)$ 
  requires {  $\mathcal{T}(R)$  }
  writes { M }
  returns { | u  $\rightarrow$   $\mathcal{T}(E)$  }
  returns { | u  $\rightarrow$   $\mathcal{T}(fE)$  }
  returns { | u  $\rightarrow$   $\mathcal{T}(Wu)$  }

let p_impl0 (t:  $\mathcal{T}(T)$ ):  $\mathcal{T}(U)$ 
  requires {  $\mathcal{T}(R)$  } requires {  $\mathcal{T}(fR)$  }
  returns { | u  $\rightarrow$   $\mathcal{T}(E)$  }
  =(
     $\mathcal{T}(\text{var } u: U; \text{var } l: L;)$ 
    assume {  $\mathcal{T}(Wg)$  } -- where of globals
    assume {  $\mathcal{T}(Wt)$  } -- where of inputs
    assume {  $\mathcal{T}(Wl)$  } -- where of locals
    assume {  $\mathcal{T}(Wu)$  } -- where of outputs
    try (  $\mathcal{T}(B)$  )
    with | Return  $\rightarrow$  assume { true } end
     $\mathcal{T}(u)$ 
  )

let p_impl0_frame (t:  $\mathcal{T}(T)$ ):  $\mathcal{T}(U)$ 
  requires {  $\mathcal{T}(R)$  } requires {  $\mathcal{T}(fR)$  }
  writes { M }
  reads { G } -- all globals
  returns { | u  $\rightarrow$  true }
  =( ... -- as in p_impl0
     $\mathcal{T}(m := m), \text{for } m \in M$ 
    assume { yes(g) }, for g  $\in G$ 
     $\mathcal{T}(u)$  )

```

Fig. 7. Translation of a Boogie procedure (left) into WhyML (right).

The specification of procedure p consists of preconditions **requires**, frame specification **modifies**, and postconditions **ensures**. A precondition is an assertion that callers of p must satisfy upon calling, and that every implementation of p can assume; **free** preconditions need not be satisfied by callers. A postcondition is an assertion that every implementation of p must satisfy upon terminating, and that every caller of p can assume; **free** postconditions need not be satisfied by implementations. Every implementation of p may only modify the global variables listed in p 's frame specification.

\mathcal{T} translates a generic procedure p as shown in Fig. 7 (right). The declaration of p determines **val** p , which defines the semantics of p for clients: the **free** precondition fR does not feature there because clients don't have to satisfy it, whereas both **free** and non-**free** postconditions are encoded as **returns** conditions. The implementation of p determines **let** p_impl0 , which triggers the verification of the implementation against its specification: both **free** and non-**free** preconditions are encoded, whereas the **free** postcondition fE does not feature there because implementations don't have to satisfy it. The body introduces **let** bindings for the local variable l and for a new local variable u which represents the returned value; these declarations are translated as discussed in Sec. 4.3. Then, a series of **assume** encode the semantics of Boogie's **where** clauses, which constrain the nondeterministic values variables can take (Wg comes from any global variables, which are visible everywhere); p 's body B is translated and wrapped inside an exception-handling block **try**, which does not do anything other than allow-

ing abrupt termination of the body's execution upon throwing a `Return` exception (see Sec. 4.7 for details). Regardless of whether the body terminates normally or exceptionally, the last computed value of `u` is returned in the last line, and checked against the postcondition in `returns`. Another implementation `let p_impl0_frame` checks the frame condition (`modifies` clause).¹² It relies on the same full precondition as `p_impl0` but has postcondition `true` since `E` has already been checked; it includes a `writes` clause and a `reads` clause. Why3 checks that a global variable is in the `writes` clause if and only if it is written by the implementation; since Boogie's `modifies` clause only expresses variables that *may* be written, `p_impl0_frame` includes an assignment of every variable in `M` to itself so that the requirement that every variable in `M` is written is vacuously satisfied. When a `writes` clause is present, Why3 also requires a `reads` clause and checks that every variable in it is written, read, or both. The translation builds a `reads` clause with all global variables `G`, and vacuously reads all of them using function `yes 'a: bool`, which identically returns `true` for any input; this makes the `reads` clause satisfied by any implementation. In all, the modular semantics of Boogie's procedure `p` is preserved.

4.7 Statements

Axioms and assertions. Boogie's `assert e`, `assume e`, and `axiom e` statements translate to `assert { T(e) }`, `assume { T(e) }`, and `axiom A: T(e)` in WhyML.

Assignments. Assignments involve variables (global or local), which become mutable references in WhyML: $\mathcal{T}(v := e) = v.\text{contents} \leftarrow \mathcal{T}(e)$. Boogie parallel assignments become simple assignments using `let` bindings of limited scope:

$$\mathcal{T}(v_1, \dots, v_m := e_1, \dots, e_m) = \left\{ \begin{array}{l} \text{let } e'_1 = \mathcal{T}(e_1), \dots, e'_m = \mathcal{T}(e_m) \text{ in} \\ \mathcal{T}(v_1 := e'_1); \dots; \mathcal{T}(v_m := e'_m) \end{array} \right. \quad (2)$$

Havoc. An abstract function `val havoc () : 'a` provides a fresh, nondeterministic¹³ value of any type `'a`. It translates Boogie's `havoc` statements following the scheme:¹⁴

$$\mathcal{T}(\text{havoc } u, v) = \mathcal{T}(u) \leftarrow \text{havoc}(); \mathcal{T}(v) \leftarrow \text{havoc}(); \text{assume } \{ \mathcal{T}(W_u) \}; \text{assume } \{ \mathcal{T}(W_v) \}$$

where `Wu` and `Wv` are the `where` clauses of `u`'s and `v`'s declarations; the generalization to an arbitrary number of variables is obvious. It is important that the `assume` statements follow all the calls to `havoc`: since `Wv` may involve `u`'s value, `havoc u, v` is not in general equivalent to `havoc u; havoc v`; the translation reflects this behavior.

Return. The behavior of Boogie's `return` statement, which determines the abrupt termination of a procedure's execution, is translated to WhyML using exception handling. An exception handling block wraps each procedure's body, as illustrated in Fig. 7, and catches an `exception Return`; thus, $\mathcal{T}(\text{return}) = \text{raise Return}$.

¹² The tool `b2w` does not currently implement frame condition checks.

¹³ <http://lists.gforge.inria.fr/pipermail/why3-club/2013-April/000615.html>

¹⁴ Alternatively, we could define $\mathcal{T}(\text{havoc } v) = \text{any } \mathcal{T}(T)$, where `T` is `v`'s type.

Jumps (branching). In addition to structured `while` loops (discussed below), Boogie provides jump statements of the form `goto ι_1, \dots, ι_n` , which nondeterministically jump to any of the locations labeled by ι_k . The translation must remove jump statements in a way that preserves verifiability; this rules out “global” approaches using a program counter [11,27], since they would require new invariants about the counter. Instead, we introduce simple heuristics that replace jumps with structured code; since the usage of jumps in Boogie programs tend to follow well-defined patterns that can be traced back to structured loops, the heuristics may be sufficient in practice.¹⁵




Fig. 8. Transformation of loops from unstructured (left) to structured (right).

Consider the control-flow graph G of a procedure body; each node N of G is a *simple block*: a linear piece of code with a label ℓ_N on the first statement, no labels anywhere else in N , and a `goto` as last statement or no `goto` statements at all; arrows connect N to the locations mentioned in N ’s `goto` statement (if N has no `goto`, we call it a *terminal node*). We apply three kinds of transformations on G exhaustively.

Sequencing: if $N \rightarrow M$ is the only arrow out of N and the only arrow into M , and $M \not\rightarrow N$, replace N and M with the single block $N; M$ with the `goto` at the end of N and label ℓ_M removed.

Choosing: if $N \rightarrow \{M_1, \dots, M_n\}$ are the only arrows out of N and the only arrows into each M_1, \dots, M_n , and every M_k , for $1 \leq k \leq n$, is a terminal node, replace N, M_1, \dots, M_n with the single block:

$$N; \text{if } (*)\{M_1\} \text{ else } \{\text{if } (*)\{M_2\} \text{ else } \{\dots \text{else } \{M_n\}\} \dots \}$$

with the `goto` at the end of N and all labels other than ℓ_N removed.¹⁶

Looping: replace the subgraph of Fig. 8 (left) with the structured loop to its right.

Conditionals. The translation of conditionals is straightforward:

$$\mathcal{T}(\text{if } (b) \text{ then } \{BT\} \text{ else } \{BE\}) = \text{if } \mathcal{T}(b) \text{ then } \{ \mathcal{T}(BT) \} \text{ else } \{ \mathcal{T}(BE) \}$$

¹⁵ \mathcal{T} ’s implementation currently does not support this translation of `goto` statements.

¹⁶ This is after Dafny’s calculational proof approach [17].

Loops. Fig. 9 shows the translation of a Boogie loop into a WhyML loop. An invariant marked as **free** can be assumed but need not be checked; correspondingly, the translation adds assumptions that ensure it holds at loop entrance and after every iteration. The exception handling block surrounding the loop in WhyML emulates the semantics of the control-flow breaking statement **break**: $\mathcal{T}(\text{break}) = \text{raise Break}$.

<pre style="margin: 0;">while (b) invariant I; free invariant fI; { B }</pre>	<pre style="margin: 0;">assume { $\mathcal{T}(fI)$ } try while $\mathcal{T}(b)$ do invariant { $\mathcal{T}(I)$ } invariant { $\mathcal{T}(fI)$ } $\mathcal{T}(B)$ assume { $\mathcal{T}(fI)$ } done; with Break \rightarrow assume { $\mathcal{T}(fI)$ } end</pre>
---	---

Fig. 9. Translation of a Boogie loop (left) into WhyML (right).

Procedure calls. The translation of procedure calls is straightforward; for Boogie procedure p in Fig. 7: $\mathcal{T}(\text{call } r := p(e)) = \mathcal{T}(t) \leftarrow p(\mathcal{T}(e))$. Since WhyML function calls translating Boogie procedures use the **val** style of declaration rather than the recursive function style (**rec**), the modular semantics of procedure calls (where the behavior is entirely determined by the specification) is correctly preserved.

Call-forall. \mathcal{T} translates call-forall statements (supported in older versions of Boogie [15]) by axiomatizing their semantics:

$$\mathcal{D}(\text{call forall Lemma}(*)) = \text{assume } (\forall t: T \bullet R(t) \implies E(t))$$

where **Lemma** is declared as **procedure** **Lemma**($t: T$) **requires** $R(t)$; **ensures** $E(t)$.

4.8 Attributes

\mathcal{T} translates *triggers* using WhyML's syntax:

$$\mathcal{T}(\forall x: X \bullet \{\text{trig}\} E(x)) = \forall x: \mathcal{T}(X) \ [\mathcal{T}(\text{trig})]. \mathcal{T}(E(x))$$

The translation discards other application-specific attributes, which have no equivalent in Why3.

4.9 Identifiers and Visibility

Boogie is more liberal than WhyML in the range of characters that are allowed in identifier names; therefore, the translation defines an injective renaming of identifiers when necessary.

Boogie allows local declarations to shadow global declarations of entities with the same name. Since WhyML does not allow shadowing, the translation introduces fresh names for local declarations when necessary to avoid name clashes with the shadowed declarations.

While the order of declarations is immaterial in Boogie, in WhyML reference must follow declaration. Thus, the translation reorders declarations to comply with WhyML's requirements; it also introduces a canonical order of declarations: types, global variables, functions, axioms, procedure declarations (**val**), procedure definitions (**let**), other declarations.

4.10 Polymorphic Maps

We now consider *polymorphic map* types, declared in Boogie as:

$$\text{type } \mathbf{pM} = \langle \alpha \rangle [T_1, \dots, T_n] \cup \quad (3)$$

where α is a vector $\alpha_1, \dots, \alpha_m$ of $m > 0$ type parameters, and some of the types T_1, \dots, T_n, \cup in \mathbf{pM} 's definition depend on α . In the next paragraph, we explain why polymorphic maps cannot be translated to WhyML directly. Instead, we replace them with several monomorphic maps based on a global analysis of the types that are actually used in the Boogie program being translated. The result of this rewrite is a Boogie program without polymorphic maps, which we can translate to Why3 following the rules we previously described. The shortcoming of this approach is that it gives up *modularity*: verification holds only for the concrete types that are used (closed-word assumption); this seems to be necessary to express Boogie's extremely liberal polymorphism without resorting to intricate "semantic" translations, which would likely fail verifiability.

Boogie vs. WhyML polymorphism. While WhyML also supports generic polymorphism, like every functional language in the ML family to which it belongs, its usage is more restrictive than Boogie's. The first difference is that *mutable* maps cannot be polymorphic in WhyML; therefore, Boogie variables of polymorphic map type require a special translation. The second difference is that, in some contexts, a variable of polymorphic map type in Boogie effectively corresponds to *multiple* maps, one for each possible concrete type, and the different maps can be combined in the same expression. Consider, for example, a **type** $\mathbf{Mix} = \langle \alpha \rangle [\alpha] \alpha$ of maps from generic type α to α ; Boogie accepts formulas such as **axiom** $(\forall m: \mathbf{Mix} \bullet m[0] = 1 \wedge m[\mathbf{true}])$ where m acts as a map over **int** in the first conjunct and as a map over **bool** in the second conjunct. WhyML, in contrast, always makes the type parameters explicit; hence, a logic variable of type map $'a \rightarrow 'a$ denotes a single map of a generic type that can only feature in expressions which do not assume anything about the concrete type that will instantiate $'a$. Note that Boogie even allows expressions that introduce inconsistencies, such as $\forall \langle \beta \rangle x: \beta, y: \mathbf{Mix} \bullet y[x] = 3 \wedge y[x] = \mathbf{true}$ (where the quantification is also type-generic), which passes typechecking but allows one to derive false.

Besides type declarations and quantifications, polymorphic maps can appear within polymorphic functions and procedures, declared as:

$$\text{function } pF\langle\alpha\rangle(x_1 : T_1, \dots, x_n : T_n) \text{ returns } (U) \quad (4)$$

$$\text{procedure } pP\langle\alpha\rangle(x_1 : T_1, \dots, x_n : T_n) \text{ returns } (u : U) \quad (5)$$

Precisely, two kinds of polymorphic maps may feature within polymorphic functions and procedures: polymorphic maps generic with respect to *explicitly declared* function or procedure parameters are similar to Why3's, and hence different from those generic with respect to implicit type parameters declared outside the function or procedure. For example, implementations of a procedure $p\langle\beta\rangle(m : \text{Mix}, n : [\beta]\beta)$ can select elements of any concrete type from m , but only elements of parametric type β from n .

Type analysis. We have seen that a Boogie polymorphic map may correspond to multiple monomorphic maps in certain contexts. The translation reifies this idea based on global type analysis: for every item (constant, program or logic variable, or formal argument) pm of polymorphic map type pM as in (3), it determines the set $types(pm)$ of all actual types pm takes in expressions or assignments, as outlined in Tab. 10.¹⁷ This in turn determines the set $types(pM)$ as the union of all sets $types(p)$ for p of type pM .

			$types(pm)$ includes $[t_1, \dots, t_n]u$ such that:
expressions	read	pm	$pm :: [t_1, \dots, t_n]u$
	select	$pm[e_1, \dots, e_n]$	$e_1 :: t_1, \dots, e_n :: t_n, pm[e_1, \dots, e_n] :: u$
	update	$pm[e_1, \dots, e_n] := f$	$e_1 :: t_1, \dots, e_n :: t_n, f :: u$
	function reference	$f(it)$	$it :: [t_1, \dots, t_n]u$, where function $f(pm : pM)$
	copy	$pm := it$	$it :: [t_1, \dots, t_n]u$
statements	assignment	$pm[e_1, \dots, e_n] := f$	$e_1 :: t_1, \dots, e_n :: t_n, f :: u$
	havoc	havoc pm	—
	procedure call in	call $p(it)$	$it :: [t_1, \dots, t_n]u$, where procedure $p(pm : pM)$
	procedure call out	call $it := p()$	$it :: [t_1, \dots, t_n]u$, where procedure $p()$ returns $(pm : pM)$

Table 10. Each occurrence of an item pm of polymorphic map type pM determines the set $types(pm)$ of actual types. ($x :: t$ denotes that x has type t .)

The types in $types(pM)$ include in general both concrete and parametric types. For example, the program of Fig. 11 (left) determines $types(m) = \{[\text{int}]\text{int}, [\beta]\beta\}$, $types(n) = \{[\text{bool}]\text{bool}\}$, and $types(M) = types(m) \cup types(n)$, where β is procedure p 's type parameter (since p is not called anywhere, that's the only known actual type of x). Let $conc(pM)$ denote the set of all *concrete* types in $types(pM)$.

Desugaring polymorphic maps. To describe how the translation replaces polymorphic maps by monomorphic maps, we introduce a pseudo-code notation that allows *tuples* (in round brackets) of program elements where normally only a single element is allowed. The semantics of this notation corresponds quite intuitively to multiple statements or

¹⁷ A parameter's actual type is ambiguous if the parameter appears in the map type's codomain but not in its domain; in this case, Boogie defaults to type **int**.

<pre> type M = ⟨α⟩ [α]α; var m: M; axiom (∀ n: M • n[true]); procedure p⟨β⟩(x: β) requires (∀ i: int • m[i] = i); modifies m; { m[x] := x; } </pre>	<pre> type (M_int, M_bool, M_a) = ([int]int, [bool]bool, [a]a); var (m_int, m_bool, m_a): (M_int, M_bool, M_a); axiom (∀ (n_int, n_bool, n_a): (M_int, M_bool, M_a) • n_bool[true]); procedure (p_int, p_bool, p_a)(x: (int, bool, a)) requires (∀ i: int • m_int[i] = i); modifies {m_int, m_bool, m_a}; { (m_int, m_bool, m_a)[x] := x } </pre>
--	--

Fig. 11. An example of how polymorphic maps (left) translate to monomorphic (right). Procedure p translates to 3 procedures p_int , p_bool , and p_a , each with argument of type `int`, `bool`, or `a`.

declarations. For example, a variable declaration `var (x, y): (int, bool)` is a shorthand for declaring variables $x: \text{int}$ and $y: \text{bool}$; a formula $(x, y) = (3, \text{true})$ is a shorthand for $x = 3 \wedge y$; and a procedure declaration using the tuple notation `procedure (p_int, p_bool)(x: (int, bool))` is a shorthand for declaring two procedures $p_int(x: \text{int})$ and $p_bool(x: \text{bool})$.

We also use the following notation: given an n -vector $\mathbf{a} = a_1, \dots, a_n$ and a type expression T parametric with respect to α , $T_{\mathbf{a}}$ denotes T with a_k substituted for α_k , for $k = 1, \dots, n$. If \mathbb{T} is a set of types obtained from the same type expression T , such as $\text{types}(\text{pM})$ with respect to pM 's definition, and id is an identifier, let (\mathbb{T}) denote \mathbb{T} as a tuple, and $(\text{id_}\mathbb{T})$ denote the tuple of identifiers $\text{id_}t$ such that T_t is the corresponding type in \mathbb{T} . In the example of Fig. 11, if $T = [\alpha]\alpha$ then $T_{\text{int}} = [\text{int}]\text{int}$, $(\text{types}(\text{m})) = ([\text{int}]\text{int}, [\beta]\beta)$, and $(\text{j_types}(\text{m})) = (\text{j_int}, \text{j_}\beta)$. Throughout, we also assume that an uninterpreted type \mathbf{a}_k is available for $k = 1, \dots, n$, that $M_{\mathbf{a}}$ denotes the type expression $[T_1, \dots, T_n] \cup$ in (3) with each α_k replaced by \mathbf{a}_k , and that $\text{conc}^+(\text{pM}) = \text{conc}(\text{pM}) \cup \{M_{\mathbf{a}}\}$.

Declarations. Type declaration (3) desugars to several type declarations:

$$\text{type } (\text{pM_conc}^+(\text{pM})) = (\text{conc}^+(\text{pM})) \quad (6)$$

The declaration of an *item* $\text{pm}: \text{pM}$, where pm can be a constant, or a program or logic variable, desugars to a declaration $(\text{pm_conc}^+(\text{pM})) : (\text{conc}^+(\text{pM}))$ of multiple items of the same kind. The declaration of a *procedure* or *function* g with an (input or output) argument $x: \text{pM}$ desugars to a declaration of multiple procedures or functions $(g_conc^+(\text{pM}))(x: (\text{conc}^+(\text{pM})))$ —multiple declarations each with one variant of x ; if g has multiple arguments of this kind, the desugaring is applied recursively to each variant. Fig. 11 (right) shows how the polymorphic map type M and each of the items m and n of type M become 3 monomorphic types and 3 items of these monomorphic types.

For every polymorphic function or procedure g with type parameters β , also consider any one of their arguments declared as $x: X$. If X is a type expression that depends on β , and there exists a map type $[V_1, \dots, V_n]V_0$ in $\text{types}(\text{pM})$ such that $X = V_k$ for some $k = 0, \dots, n$, then g becomes $(g_V_k)(x: (V_k))$ —corresponding to multiple g 's each with one argument, where $V_k = \{\bar{V}_k \mid [\bar{V}_1, \dots, \bar{V}_n]\bar{V}_0 \in \text{conc}^+(\text{pM})\}$ is

the set of all concrete types that instantiate the k th type component. This transformation enables assigning arguments to polymorphic maps inside polymorphic functions or procedures that have become monomorphic. Fig. 11 (right) shows how argument $x: \beta$ becomes an argument of concrete type `int`, `bool`, or `a`, since $[\beta]\beta \in \text{types}(\mathcal{M})$. (As procedure p does not use β elsewhere, we drop it from the signature.)

Expressions. Every occurrence—in expressions, as l-values of assignments, and as targets of `havoc` statements—of an item w of polymorphic type W whose declaration has been modified to remove polymorphic map types is replaced by one or more of the newly introduced monomorphic types as follows. If w ’s actual type within its context is a *concrete* type C , then we replace w with $w.c$ such that $W_c = C$; otherwise, w ’s actual type is a *parametric* type, and we replace w with the tuple $(w.X)$, including all variants of w that have been introduced. In Fig. 11 (right), $n[\text{true}]$ rewrites to just $n.\text{bool}[\text{true}]$ since the concrete type is `bool`; the assignment in p ’s body, whose actual type is parametric with respect to β , becomes an assignment involving each of the three variants of m corresponding to the three variants of p that have been introduced.

5 Implementation and Experiments

5.1 Implementation

We implemented the translation \mathcal{T} described in Sec. 4 as a command-line tool `b2w` implemented in Java 8. `b2w` works as a staged filter: 1) it parses and typechecks the input Boogie program, and creates a Boogie AST (abstract syntax tree); 2) it desugars the Boogie AST according to \mathcal{D} ; 3) it transforms the Boogie AST into a WhyML AST according to \mathcal{E} ; 4) it outputs the WhyML AST in the form of code.

Stage 1) relies on Schäfer’s parsing and typechecking library `Boogieamp`¹⁸, which we modified to support access using the visitor pattern, AST in-place modifications, and the latest syntax of Boogie (e.g., for integer vs. real division¹⁹). Stages 2) and 3) are implemented by multiple AST visitors, each taking care of a particular aspect of the translation, in the style of [27]; the overhead of traversing the AST multiple times is negligible and improves modularity: handling a new construct (for example, in future versions of Boogie) or changing the translation of one feature only requires adding or modifying one feature-specific visitor class. A similar technique is also advocated in [22].

5.2 Experiments

The goal of the experiments is ascertaining that `b2w` can translate realistic Boogie programs producing WhyML programs that can be verified taking advantage of Why3’s multiple back-end support. The experiments are limited to fully-automated verification, and hence do not evaluate other possible practical benefits of translating programs to WhyML such as support for interactive provers and executability for testing purposes.

¹⁸ <https://github.com/martinschaef/boogieamp>

¹⁹ <http://boogie.codeplex.com/discussions/397357>

Programs. The experiments target a total of 194 Boogie programs from three groups according to their origin: group NAT (native) includes 29 programs that encode algorithmic verification problems directly in Boogie (as opposed to translating from a higher-level language); group OBJ (object-oriented) includes 6 programs that are based on a heap-based memory model; group TES (tests) includes 159 programs from Boogie’s test suite. Tab. 12 summarizes the sizes of the programs in each group.

GROUP	#	LOC BOOGIE				LOC WHYML			
		m	μ	M	Σ	m	μ	M	Σ
NAT	29	20	73	253	2110	62	128	318	3716
OBJ	6	44	146	385	878	90	208	446	1245
TES	159	3	21	155	3272	36	64	290	10180
Total:	194	3	34	385	6260	36	106	446	15141

Table 12. A summary of the Boogie programs used in the experiments, and their translation to WhyML using b2w. For each program GROUP, the table reports how many programs it includes (#), the minimum m , mean μ , maximum M , and total Σ length in non-comment non-blank lines of code (LOC) of those BOOGIE programs and of their WHYML translations.

The programs in NAT, which we developed in previous work [10,9], include several standard algorithms such as sorting and array rotation. The programs in OBJ include 2 simple examples in Java and 1 in Eiffel, encoded in Boogie by Joogie [1] and AutoProof [28] (we manually simplified AutoProof’s translation to avoid features b2w doesn’t support), and 3 algorithmic examples adapted from NAT to use a global heap in the style of object-oriented programs. Among the 515 programs that make up Boogie’s test suite²⁰ we retained in TES those that mainly exercise features supported by b2w. This meant excluding several groups of tests that exercise special options (Houdini, assertion inference, special Z3 encodings and directives, etc.), unsupported language features (bitvectors, gotos, etc.), and the correctness of typechecking (b2w assumes well-formed Boogie input). It also meant excluding 4 programs that triggered Boogie errors (a Boogie *error* means here a problem with the input such as a typechecking or parsing error due to a feature not activated; it is not a verification error, which just denotes a failed verification attempt and is fair game for evaluating the translation); and another 35 programs that b2w failed to translate because of unsupported features that we identified a posteriori.

Setup. Each experiment targets one Boogie program b : it runs Boogie with command `boogie b` and a timeout of 180 seconds; it runs b2w to translate b to w in WhyML; for each SMT solver p among Alt-Ergo, CVC3, CVC4, and Z3, it runs Why3 with command `why3 prove -P p w`, also with a timeout of 180 seconds.²¹ For each run we collected the wall-clock running time, the total number of verification goals, and how many of such goals the tool verified successfully.²²

²⁰ <https://github.com/boogie-org/boogie/tree/master/Test>

²¹ The timeouts were enforced using the Unix command `timeout`. We also set had a 20-second timeout per procedure (option `/timeLimit` in Boogie) or goal (option `-T` in Why3).

²² The number of verification goals of each program is the same in Boogie and Why3: the number of procedure implementations.

All the experiments ran on a Ubuntu 14.04 LTS GNU/Linux box with 8-core Intel i7-4790 CPU at 3.6 GHz and 16 GB of RAM, with the following tools: Alt-Ergo 0.99.1, CVC3 2.4.1, CVC4 1.4, Z3 4.3.2, Mono 4.2.2, OCaml 4.02.3, Boogie 2.3.0.61016, and Why3 0.86.2. To account for noise, we repeated each verification three times and report the mean value of the 95th percentile of the running times.

GROUP	#	B = W	B > W	B < W	0=0	50=50	100=100	SPURIOUS
NAT	29	19	10	0	1	0	18	0
OBJ	6	5	0	1	1	2	2	0
TES	159	137	21	1	71	21	45	0
Total:	194	161	31	2	73	23	65	0

Table 13. A summary of how Boogie performs in comparison with Why3. For each program GROUP, the table reports how many programs it includes (#), for how many of the programs Boogie verifies as many goals ($B = W$), more goals ($B > W$), or fewer goals ($B < W$) than Why3 with any of the SMT solvers; for how many of the programs both Boogie and Why3 verify none ($0=0$), some but not all ($50=50$), or all ($100=100$) of the goals; the last column (SPURIOUS) indicates that b2w’s translation never introduces spurious goals that are proved by Why3 (that is, if Boogie’s input has zero goals, so does WhyML’s translation).

Results. Tab. 13 shows a summary of the results where we compare Why3’s best performance, with any one of the four SMT solvers, against Boogie’s. The most significant result is that the WhyML translation produced by b2w behaves like the Boogie original in 83% (161, $B=W$) of the experiments. This means that Boogie may fail to verify all goals (column $0=0$), verify some goals and fail on others (column $50=50$), or verify all goals (column $100=100$); in each case, Why3 consistently verifies the same goals on b2w’s translation. Indeed, many programs in TES are tests that are supposed to fail verification; hence, the correct behavior of the translation is to fail as well. We also checked the failures of programs in NAT and OBJ to ascertain that b2w’s translation preserves correctness. Tab. 13 does not show this, but we also found another 2 programs in NAT (*inv_survey/bst* and *rotation/rotation_reverse*) where Why3 proves the same goals as Boogie only by combining the results of multiple SMT solvers.

Boogie verifies more goals than Why3 in 16% (31, $B > W$) of the experiments, where it is more effective because of better features (default triggers, invariant inference, SMT encoding) or simply because of some language features that are not fully supported by b2w (examples are Z3-specific annotations, which b2w simply drops, and *goto*, which b2w encodes as *assert false* to ensure soundness). In 1% (2, $B < W$) of the experiments, Why3 even verifies more goals than Boogie. One program in OBJ (*rotation_by_copy*) is a genuine example where Why3’s Z3 encoding is more effective than Boogie’s²³; the one program in TES (*test2/Quantifiers*) should instead be considered spurious, as it deploys some trigger specifications that are Boogie-specific (negated triggers) or interact in a different way with the default triggers. (Procedures S, U0, and U1 use regular triggers whose translation to Why3 yields a different behavior, probably because of Why3’s default triggers differ from Boogie’s; procedures W and X2

²³ However, Boogie also succeeds given a longer timeout than the one used in the experiments.

use negated triggers, which b2w ignore.) As this was the only program in our experiments that introduced clearly spurious behavior, the experiments provide convincing evidence that b2w’s translation preserves correctness and verifiability to a large degree.

GROUP	#	Z3 BOOGIE			ALT-ERGO WHY3			CVC3 WHY3			CVC4 WHY3			Z3 WHY3		
		OUTCOME		TIME	OUTCOME		TIME	OUTCOME		TIME	OUTCOME		TIME	OUTCOME		TIME
		μ	\forall	\exists	μ	\forall	\exists	μ	\forall	\exists	μ	\forall	\exists	μ	\forall	\exists
NAT	29	93	25	1	0.4	12	0	61	14	6	20.6	598	0	28	1	12
OBJ	6	52	2	2	3.9	23	0	46	1	2	30.1	181	0	46	1	2
TES	159	45	55	71	0.3	53	0	37	45	85	25.8	4096	1	33	39	91
Total:	194	60	82	74	0.7	88	0	53	60	93	22.6	4875	1	30	41	105

Table 14. For each program GROUP the table reports how many programs it includes (#) and, for both Boogie and Why3 for each choice of SMT solver among ALT-ERGO, CVC3, and Z3: the mean percentage of goals verified in each program (OUTCOME μ), how many programs were completely verified (OUTCOME \forall), and how many were not verified at all (OUTCOME \exists), the mean μ and total Σ verification TIME in seconds (including time outs), and how many programs timed out.

Tab. 14 provides data about the experiments’ running times, and differentiates the performance of each SMT solver with Why3. Z3 is the most effective SMT solver in terms of programs it could completely verify (columns \forall), followed by Alt-Ergo. While CVC3 is generally the least effective, it has the advantage of returning very quickly (only 0.2 seconds of average running time), even more quickly than Z3 in Boogie. CVC4 falls somewhere in the middle, in terms both of effectiveness and of running time. Boogie’s responsiveness remains excellent if balanced against its effectiveness; a better time-effectiveness of Why3 with Alt-Ergo and Z3 could be achieved by setting tight per-goal timeouts (in most cases, verification attempts that last longer than a few seconds do not eventually succeed).

6 Discussion

The current implementation of the translation \mathcal{T} has some limitations that somewhat restrict its applicability. As we already mentioned in the paper, some features of the Boogie language are not supported (bitvectors, gotos), or only partially supported (polymorphic mappings); and frame specifications are assumed. All of these are, however, limitations of the current prototype implementation only, and we see no fundamental hurdles to extending b2w along the lines of the definition of \mathcal{T} in Sec. 4.

Since b2w also takes great care to confine the effect of translating Boogie programs that include unsupported features, and to fail when it cannot produce a correct translation, it still largely preserves *correctness* (soundness, in particular). For example, a **goto** statement is rendered as **assert false**; therefore, the translated program verifies only if the **goto** is never executed in the original program, which ensures soundness. On the other hand, our experiments also demonstrate that the translation \mathcal{T} , as implemented by b2w, largely meets the other goal of preserving *verifiability*: even if the experimental subjects all are idiomatic Boogie programs written independent of the translation effort, 83% of the translated programs behave in Why3 as they do in Boogie.

In future work, we will address the features of Boogie that are still not satisfactorily supported by b2w. We will also devise strategies to take advantage of Why3’s multi-prover support. Other possible directions include formalizing the translation to prove that it preserves correctness; and devising a reverse translation from WhyML to Boogie.

References

1. S. Arlt and M. Schäf. Joogie: Infeasible code detection for Java. In *Proceedings of CAV*, volume 7358 of *LNCS*, pages 767–773. Springer, 2012.
2. M. Barnett, B. E. Chang, R. DeLine, B. Jacobs, and K. R. M. Leino. Boogie: A modular reusable verifier for object-oriented programs. In *Proceedings of FMCO*, volume 4111 of *LNCS*, pages 364–387. Springer, 2006.
3. M. Barnett, M. Fähndrich, K. R. M. Leino, P. Müller, W. Schulte, and H. Venter. Specification and verification: the Spec# experience. *Commun. ACM*, 54(6):81–91, 2011.
4. S. Böhme, K. R. M. Leino, and B. Wolff. HOL-Boogie – an interactive prover for the Boogie program-verifier. In *Proceedings of TPHOLs*, volume 5170, pages 150–166. Springer, 2008.
5. Z. Cheng, R. Monahan, and J. F. Power. A sound execution semantics for ATL via translation validation – research paper. In *Proceedings of ICMT*, volume 9152 of *LNCS*, pages 133–148. Springer, 2015.
6. M. Clochard, J. Filliâtre, C. Marché, and A. Paskevich. Formalizing semantics with an automatic program verifier. In *Proceedings of VSTTE*, volume 8471 of *LNCS*, pages 37–51, 2014.
7. L. M. de Moura and N. Bjørner. Z3: An efficient SMT solver. In *Proceedings of TACAS*, pages 337–340, 2008.
8. J. Filliâtre and A. Paskevich. Why3 – where programs meet provers. In *Proceedings of ESOP*, volume 7792 of *LNCS*, pages 125–128. Springer, 2013.
9. C. A. Furia. Rotation of sequences: Algorithms and proofs. <http://arxiv.org/abs/1406.5453>, June 2014.
10. C. A. Furia, B. Meyer, and S. Velder. Loop invariants: Analysis, classification, and examples. *ACM Computing Surveys*, 46(3):Article 34, 2014.
11. D. Harel. On folk theorems. *Commun. ACM*, 23(7):379–389, 1980.
12. S. Heule, I. T. Kassios, P. Müller, and A. J. Summers. Verification condition generation for permission logics with abstract predicates and abstraction functions. In *Proceedings of ECOOP*, volume 7920 of *LNCS*, pages 451–476. Springer, 2013.
13. G. Klein, J. Andronick, K. Elphinstone, G. Heiser, D. Cock, P. Derrin, D. Elkaduwe, K. Engelhardt, R. Kolanski, M. Norrish, T. Sewell, H. Tuch, and S. Winwood. seL4: formal verification of an operating-system kernel. *Commun. ACM*, 53(6):107–115, 2010.
14. R. Kumar, M. O. Myreen, M. Norrish, and S. Owens. CakeML: a verified implementation of ML. In *Proceedings of POPL*, pages 179–192. ACM, 2014.
15. K. R. M. Leino. This is Boogie 2, 2008. <http://goo.gl/QsH6g>.
16. K. R. M. Leino. Developing verified programs with Dafny. In *Proceedings of ICSE*, pages 1488–1490. ACM, 2013.
17. K. R. M. Leino and N. Polikarpova. Verified calculations. In *Proceedings of VSTTE*, pages 170–190, 2013.
18. X. Leroy. Formal verification of a realistic compiler. *Commun. ACM*, 52(7):107–115, 2009.
19. S. Liew, C. Cadar, and A. Donaldson. Symbooglix: A symbolic execution engine for Boogie programs. In *Proceedings of ICST*. IEEE, 2016.
20. T. Mens and P. Van Gorp. A taxonomy of model transformation. *Electr. Notes Theor. Comput. Sci.*, 152:125–142, 2006.

21. N. Polikarpova, C. A. Furia, and S. West. To run what no one has run before: Executing an intermediate verification language. In *Proceedings of RV*, volume 8174 of *LNCs*, pages 251–268. Springer, 2013.
22. D. Sarkar, O. Waddell, and R. K. Dybvig. Educational pearl: A nanopass framework for compiler education. *J. Funct. Program.*, 15(5):653–667, 2005.
23. P. H. Schmitt, M. Ulbrich, and B. Weiß. Dynamic frames in Java dynamic logic. In *Proceedings of FoVeOOS*, volume 6528 of *LNCs*, pages 138–152. Springer, 2011.
24. L. Segal and P. Chalin. A comparison of intermediate verification languages: Boogie and Sireum/Pilar. In *Proceedings of VSTTE*, volume 7152 of *LNCs*, pages 130–145. Springer, 2012.
25. P. Stevens. A landscape of bidirectional model transformations. In *Proceedings of GTTSE*, volume 5235 of *LNCs*, pages 408–424. Springer, 2008.
26. M. Trudel, C. A. Furia, M. Nordio, and B. Meyer. Really automatic scalable object-oriented reengineering. In *Proceedings of ECOOP*, volume 7920 of *LNCs*, pages 477–501. Springer, 2013.
27. M. Trudel, C. A. Furia, M. Nordio, B. Meyer, and M. Oriol. C to O-O translation: Beyond the easy stuff. In *Proceedings of WCRE*, pages 19–28. IEEE Computer Society, October 2012.
28. J. Tschannen, C. A. Furia, M. Nordio, and N. Polikarpova. AutoProof: Auto-active functional verification of object-oriented programs. In *Proceedings of TACAS*, volume 9035 of *LNCs*, pages 566–580. Springer, 2015.

NAME	BOOGIE						WHY3					
	Z3			ALT-ERGO			CVC3			CVC4		
	LOC	% V.	T	LOC	% V.	T	% V.	T	% V.	T	% V.	T
inv_survey/array_partitioning_v1	42	100	0.5	100	50	20.1	50	0.2	50	21.3	50	20.3
inv_survey/array_partitioning_v2	53	100	0.4	125	100	0.3	50	0.1	50	21.3	100	0.1
inv_survey/array_stack_reversal	125	100	0.5	204	100	0.3	86	0.2	71	42.5	86	20.3
inv_survey/bst	153	100	0.4	258	50	40.5	50	0.2	75	21.4	75	20.3
inv_survey/bubble_sort_basic	49	100	0.5	113	100	0.2	50	0.1	50	21.3	100	0.2
inv_survey/bubble_sort_improved	53	100	0.5	118	100	0.6	50	0.2	50	21.3	100	0.1
inv_survey/comb_sort	56	100	0.4	124	100	0.3	50	0.2	50	21.3	100	0.2
inv_survey/dutch_flag	63	100	0.5	133	50	20.1	50	0.2	50	21.3	100	0.1
inv_survey/insertion_sort	47	100	0.6	100	0	20.1	0	0.1	0	21.3	100	0.2
inv_survey/knapsack	50	100	0.3	97	100	7.5	0	0.1	0	21.2	100	0.1
inv_survey/Levenshtein_distance	43	100	0.3	91	100	0.4	0	0.1	100	0.9	100	0.1
inv_survey/max_of_array_v1	20	100	0.4	66	100	0.1	0	0.1	0	21.2	100	0.1
inv_survey/max_of_array_v2	20	100	0.4	66	100	0.2	0	0.1	0	21.2	100	0.1
inv_survey/partition	63	100	0.4	137	100	0.8	50	0.1	50	21.3	100	0.1
inv_survey/plateau	43	100	0.5	84	0	20.1	0	0.1	0	21.2	0	20.3
inv_survey/reverse	68	100	0.4	131	100	0.3	0	0.1	0	21.3	100	0.1
inv_survey/selection_sort	72	100	0.4	160	100	4.9	33	0.2	33	42.4	100	0.2
inv_survey/sequential_search_v1	28	100	0.4	72	0	20.1	0	0.1	0	21.2	0	20.2
inv_survey/sequential_search_v2	23	100	0.4	70	0	20.1	0	0.1	0	21.2	0	20.2
inv_survey/sum_of_array	21	100	0.3	62	100	0.1	100	0.1	0	21.2	100	0.1
inv_survey/welfare_crook	44	100	0.3	86	100	0.1	0	0.1	100	0.8	0	20.2
rotation/rotation_copy	57	100	0.4	128	33	40.1	33	0.2	33	42.4	67	20.3
rotation/rotation_copy_plain	41	100	0.3	80	0	20.1	0	0.1	0	21.2	100	0.2
rotation/rotation_reverse	201	90	0.4	318	40	120.5	10	0.4	50	106.6	80	40.6
rotation/rotation_swap-1_3	48	0	0.3	88	0	20.1	0	0.1	0	21.2	0	20.2
rotation/rotation_swap-2_3	175	60	0.4	201	20	80.2	20	0.2	20	84.7	40	60.6
rotation/rotation_swap-3_3	47	100	0.3	96	67	20.1	67	0.2	67	21.3	100	0.2
rotation/rotation_swap_iterative-1_2	152	100	0.5	184	33	40.2	33	0.2	33	42.4	67	20.3
rotation/rotation_swap_iterative-2_2	253	60	0.4	224	20	80.2	20	0.3	20	84.7	40	60.6
oo/autoproof_account	385	0	0.4	446	0	80.2	0	0.4	0	84.8	0	80.8
oo/binary_search	68	100	0.4	158	67	20.1	67	0.2	100	0.2	100	0.2
oo/joogie_examples	187	60	0.5	277	60	40.2	60	0.2	60	42.4	60	40.4
oo/joogie_helloWorld	142	50	0.5	175	50	20.1	50	0.1	50	21.3	50	20.3
oo/linked_list_max	44	100	0.4	90	100	0.2	100	0.1	100	0.2	100	0.1
oo/rotation_by_copy	52	0	21.0	99	0	20.1	0	0.1	0	21.3	100	0.3

Table 15. Results for the programs in groups NAT (above the horizontal line) and OBJ (below it) in the experiments. For each program (NAME) the Boogie program length in non-comment non-empty lines of code (LOC) and the length of its WHY3 translation; and, for both Boogie and Why3, for each choice of SMT solver among ALT-ERGO, CVC3, and Z3: the percentage of goals verified in each program (% V.) and the verification time (T) in seconds (with a timeout of 180 seconds).

NAME	BOOGIE			WHY3								
	Z3			ALT-ERGO			CVC3			CVC4		
	LOC	% V.	T	LOC	% V.	T	% V.	T	% V.	T	% V.	T
doomed/doomdebug	36	0	0.3	86	0	40.1	0	0.1	0	42.4	0	40.4
doomed/doomed	73	43	0.3	185	43	80.2	43	0.1	43	84.7	43	80.7
doomed/notdoomed	43	50	0.3	107	50	40.1	50	0.1	50	42.4	50	40.4
doomed/smoke0	61	67	0.3	148	67	40.1	67	0.2	67	42.4	67	40.4
lock/Lock	86	100	0.4	163	67	20.1	67	0.1	67	21.3	67	20.2
lock/LockIncorrect	34	0	0.3	64	0	20.1	0	0.1	0	21.2	0	20.2
smoke/smoke0	41	100	0.3	108	100	0.1	100	0.1	100	0.2	100	0.1
snapshots/Snapshots0.v0	16	0	0.3	72	0	80.1	0	0.1	0	84.6	0	80.6
snapshots/Snapshots0.v1	16	50	0.3	72	50	40.1	50	0.1	50	42.4	50	40.3
snapshots/Snapshots0.v2	12	67	0.3	60	67	20.1	67	0.1	67	21.3	67	20.3
snapshots/Snapshots1.v0	10	50	0.3	48	50	20.1	50	0.1	50	21.3	50	20.2
snapshots/Snapshots1.v1	10	50	0.3	48	50	20.1	50	0.1	50	21.2	50	20.3
snapshots/Snapshots1.v2	11	50	0.3	50	50	20.1	50	0.1	50	21.2	50	20.2
snapshots/Snapshots10.v0	14	100	0.3	48	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots10.v1	14	100	0.3	48	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots11.v0	10	0	0.3	43	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots11.v1	10	0	0.3	43	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots12.v0	12	100	0.4	41	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots12.v1	12	0	0.3	41	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots13.v0	16	100	0.3	43	100	0.1	100	0.0	100	0.1	100	0.1
snapshots/Snapshots13.v1	12	0	0.4	41	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots14.v0	16	100	0.3	43	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots14.v1	16	0	0.3	43	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots15.v0	11	100	0.3	42	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots15.v1	11	0	0.3	42	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots16.v0	11	100	0.3	40	100	0.1	0	0.1	100	0.1	100	0.1
snapshots/Snapshots16.v1	11	0	0.3	40	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots17.v0	22	100	0.3	61	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots17.v1	22	0	0.3	61	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots18.v0	18	100	0.3	53	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots18.v1	18	0	0.3	53	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots19.v0	8	0	0.3	39	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots19.v1	8	0	0.3	39	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots2.v0	9	100	0.3	38	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots2.v1	9	100	0.3	38	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots2.v2	10	100	0.3	40	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots2.v3	10	100	0.3	40	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots2.v4	10	100	0.3	40	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots2.v5	11	100	0.3	42	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots20.v0	16	0	0.3	44	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots20.v1	16	0	0.3	44	0	20.1	0	0.1	0	21.3	0	20.2
snapshots/Snapshots21.v0	13	0	0.3	41	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots21.v1	13	0	0.3	41	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots22.v0	13	0	0.3	41	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots22.v1	13	100	0.3	41	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots23.v0	17	50	0.3	52	50	20.1	50	0.1	50	21.2	50	20.2
snapshots/Snapshots23.v1	18	50	0.3	53	50	20.1	50	0.1	50	21.3	50	20.2
snapshots/Snapshots23.v2	17	50	0.3	52	50	20.1	50	0.1	50	21.2	50	20.2
snapshots/Snapshots24.v0	23	0	0.3	51	0	20.1	0	0.1	0	21.3	0	20.2
snapshots/Snapshots24.v1	23	0	0.3	51	0	20.1	0	0.1	0	21.3	0	20.2
snapshots/Snapshots25.v0	11	0	0.3	47	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots25.v1	11	0	0.3	47	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots26.v0	11	0	0.3	47	0	20.1	0	0.1	0	21.3	0	20.2
snapshots/Snapshots26.v1	12	0	0.3	48	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots27.v0	11	0	0.3	47	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots27.v1	13	0	0.3	51	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots28.v0	11	100	0.3	48	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots28.v1	12	0	0.4	48	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots29.v0	11	100	0.3	47	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots29.v1	11	0	0.3	47	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots3.v0	13	100	0.3	41	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots3.v1	13	0	0.3	41	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots30.v0	11	0	0.3	42	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots30.v1	12	0	0.3	43	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots31.v0	12	100	0.3	44	100	0.1	100	0.1	100	0.1	100	0.1

NAME	BOOGIE						WHY3					
	Z3			ALT-ERGO			CVC3			CVC4		
	LOC	% V.	T	LOC	% V.	T	% V.	T	% V.	T	% V.	T
snapshots/Snapshots31.v1	11	0	0.4	43	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots32.v0	12	100	0.4	44	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots32.v1	9	0	0.3	41	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots33.v0	12	100	0.3	44	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots33.v1	6	100	0.4	37	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots34.v0	6	100	0.3	38	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots34.v1	5	0	0.3	36	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots35.v0	6	100	0.3	38	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots35.v1	5	0	0.3	36	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots36.v0	11	100	0.4	44	100	0.1	0	0.1	100	0.1	100	0.1
snapshots/Snapshots36.v1	11	0	0.3	44	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots37.v0	7	100	0.3	42	100	0.1	0	0.1	100	0.1	100	0.1
snapshots/Snapshots37.v1	7	0	0.3	42	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots38.v0	10	100	0.4	43	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots38.v1	11	0	0.3	44	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots38.v2	11	100	0.3	44	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots39.v0	10	100	0.3	43	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots39.v1	11	0	0.4	44	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots39.v2	11	100	0.3	44	100	0.1	0	0.1	100	0.1	100	0.1
snapshots/Snapshots4.v0	23	100	0.3	64	100	0.1	100	0.1	100	0.1	100	0.2
snapshots/Snapshots4.v1	27	50	0.3	76	50	40.1	50	0.1	50	42.4	50	40.4
snapshots/Snapshots40.v0	11	0	0.3	44	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots40.v1	12	0	0.3	45	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots40.v2	12	0	0.3	45	0	20.0	0	0.1	0	21.2	0	20.2
snapshots/Snapshots41.v0	31	40	0.3	99	40	60.2	40	0.1	40	63.5	40	60.5
snapshots/Snapshots41.v1	31	40	0.3	100	40	60.1	40	0.2	40	63.5	40	60.5
snapshots/Snapshots5.v0	9	100	0.3	39	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots5.v1	9	0	0.3	39	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots6.v0	12	100	0.3	43	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots6.v1	12	0	0.3	43	0	20.1	0	0.1	0	21.2	0	20.2
snapshots/Snapshots7.v0	14	100	0.3	45	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots7.v1	14	100	0.3	45	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots8.v0	11	100	0.3	45	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots8.v1	11	100	0.3	45	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots9.v0	13	100	0.4	47	100	0.1	100	0.1	100	0.1	100	0.1
snapshots/Snapshots9.v1	11	100	0.3	45	100	0.1	100	0.1	100	0.1	100	0.1
test13/ErrorTraceTestLoopInvViolationBPL	19	0	0.3	86	0	60.1	0	0.1	0	63.5	0	60.5
test15/CaptureState	23	0	0.3	62	0	20.1	0	0.1	0	21.2	0	20.2
test15/InterpretedFunctionTests	15	0	0.3	66	0	60.1	0	0.1	0	63.5	0	60.5
test15/IntInModel	3	0	0.3	36	0	20.1	0	0.1	0	21.2	0	20.2
test15/ModelTest	10	0	0.3	49	0	20.1	0	0.1	0	21.2	0	20.2
test15/NullInModel	5	0	0.3	39	0	20.1	0	0.1	0	21.2	0	20.2
test16/LoopUnroll	63	0	0.3	124	0	60.1	0	0.1	0	63.5	0	60.5
test17/contractinfer	21	0	0.3	68	0	40.1	0	0.1	0	42.4	0	40.3
test2/AssertVerifiedUnder0	26	50	0.3	100	0	120.2	0	0.2	0	126.9	0	120.9
test2/AssumeEnsures	53	57	0.3	124	57	60.2	57	0.2	57	63.6	57	60.5
test2/AssumptionVariables0	44	50	0.3	137	0	120.2	0	0.1	0	126.9	0	120.9
test2/Axioms	24	67	0.3	73	67	20.1	67	0.1	67	21.3	67	20.2
test2/B	65	100	0.3	112	0	80.1	0	0.1	0	84.6	0	80.7
test2/Call	49	40	0.3	117	20	80.2	20	0.1	20	84.6	20	80.6
test2/ContractEvaluationOrder	26	25	0.3	101	25	60.1	25	0.1	25	63.5	25	60.5
test2/CutBackEdge	35	20	0.3	96	0	100.2	0	0.1	0	105.8	0	100.7
test2/Ensures	61	50	0.5	168	50	100.2	50	0.3	50	105.9	50	100.8
test2/False	14	100	0.3	54	100	0.1	100	0.1	100	0.1	100	0.1
test2/FormulaTerm2	36	50	0.3	104	50	40.1	50	0.2	50	42.4	50	40.4
test2/FreeCall	59	64	0.5	185	27	160.2	27	0.2	27	169.2	27	161.2
test2/Implies	28	0	0.3	97	0	100.1	0	0.1	0	105.7	0	100.7
test2/InvariantVerifiedUnder0	42	17	0.3	146	0	120.2	0	0.2	0	126.9	0	120.9
test2/LoopInvAssume	15	0	0.3	44	0	20.1	0	0.1	0	21.2	0	20.2
test2/Passification	155	64	0.5	290	18	180.0	18	0.3	18	180.0	18	180.0
test2/Quantifiers	122	57	0.4	254	86	40.3	64	0.3	79	63.8	93	20.3
test2/SelectiveChecking	31	25	0.3	121	0	80.1	0	0.1	0	84.6	0	80.6
test2/sk_hack	17	100	0.3	44	0	20.1	0	0.1	0	21.2	0	20.2
test2/Timeouts0	71	0	1.3	156	0	60.1	0	0.2	0	63.5	0	60.5
test2/TypeEncodingM	19	0	0.3	60	0	20.1	0	0.1	0	21.2	0	20.2

NAME	BOOGIE			WHY3								
	Z3			ALT-ERGO			CVC3			CVC4		
	LOC	% V.	T	LOC	% V.	T	% V.	T	% V.	T	% V.	T
test21/BooleanQuantification2	9	0	0.3	46	0	20.1	0	0.0	0	21.2	0	20.2
test21/Boxing	15	0	0.5	49	0	20.1	0	0.1	0	21.2	0	20.2
test21/Casts	7	0	0.3	48	0	20.1	0	0.1	0	21.2	0	20.2
test21/Colors	13	0	0.3	62	0	40.1	0	0.1	0	42.4	0	40.4
test21/DisjointDomains	21	0	0.4	81	0	60.1	0	0.2	0	63.5	0	60.5
test21/EmptySetBug	18	0	0.4	54	0	20.1	0	0.1	0	21.2	0	20.2
test21/FunAxioms	24	50	0.5	81	50	20.1	50	0.1	50	21.3	50	20.2
test21/FunAxioms2	13	0	0.3	49	0	20.1	0	0.1	0	21.2	0	20.2
test21/InterestingExamples3	17	67	0.4	71	33	40.1	33	0.2	33	42.4	33	40.4
test21/InterestingExamples5	9	100	0.3	45	100	0.1	0	0.1	100	0.1	100	0.1
test21/Keywords	5	100	0.3	38	100	0.1	100	0.1	100	0.1	100	0.1
test21/LargeLiterals0	12	0	0.3	46	0	20.1	0	0.1	0	21.2	0	20.2
test21/LetSorting	11	100	0.3	43	0	20.1	0	0.1	0	21.2	0	20.2
test21/Maps2	14	100	0.3	52	100	0.1	0	0.1	100	0.2	0	20.2
test21/Orderings	13	50	0.4	59	0	40.1	0	0.1	0	42.4	0	40.4
test21/Orderings2	11	0	0.4	49	0	20.1	0	0.1	0	21.2	0	20.2
test21/Orderings3	22	0	0.4	78	0	40.1	0	0.1	0	42.4	0	40.4
test21/Orderings4	7	0	0.4	47	0	20.1	0	0.1	0	21.2	0	20.2
test21/PolyList	35	0	0.4	91	0	40.1	0	0.2	0	42.4	0	40.4
test21/Triggers0	34	50	0.4	92	50	20.1	50	0.1	0	42.4	50	20.2
test21/Triggers1	12	0	0.4	50	0	20.1	0	0.1	0	21.2	0	20.2
test7/MultipleErrors	14	0	0.3	42	0	20.1	0	0.1	0	21.2	0	20.2
test7/NestedVC	20	50	0.3	61	0	40.1	0	0.1	0	42.4	0	40.3
test7/UnreachableBlocks	34	100	0.3	79	50	40.1	50	0.1	50	42.4	50	40.4
textbook/Bubble	47	100	0.4	110	0	20.1	0	0.1	0	21.3	0	20.2
textbook/DutchFlag	47	100	0.3	92	0	20.1	0	0.1	0	21.2	0	20.2
textbook/Find	27	100	0.3	72	50	20.1	50	0.1	50	21.3	50	20.3
textbook/McCarthy-91	11	100	0.3	47	100	0.1	100	0.1	100	0.1	100	0.1
textbook/TuringFactorial	27	100	0.3	81	0	20.1	0	0.1	0	21.2	0	20.2

Table 16: Results for the programs in group TES in the experiments. The measures are the same as in Tab. 15.